PNAS Nexus, 2021, pp. 1–12

doi: DOI HERE
Advance Access Publication Date: Day Month Year
Paper

PAPER

Emergent structures of attention on social media are driven by amplification and triad transitivity

Alyssa Hasegawa Smith,[a,*] Jon Green,[b] Brooke Foucault Welles[a] and David Lazer[a]

[a]Network Science Institute, Northeastern University, 177 Huntington Avenue, 02115, Massachusetts, USA and [b]Political Science, Duke University, Gross Hall, 140 Science Dr, 27708, North Carolina, USA
*To whom correspondence should be addressed: smith.alyss@northeastern.edu
FOR PUBLISHER ONLY Received on Date Month Year; accepted on Date Month Year

Abstract

As they evolve, social networks tend to form transitive triads more often than random chance and structural constraints would suggest. However, the mechanisms by which triads in these networks *become* transitive are largely unexplored. We leverage a unique combination of data and methods to demonstrate a causal link between amplification and triad transitivity in a directed social network. Additionally, we develop the concept of the "attention broker," an extension of the previously theorized *tertius iungens* (or "third who joins"). We use a novel technique to identify time-bounded Twitter/X following events, and then use difference-in-differences to show that attention brokers cause triad transitivity by amplifying content. Attention brokers intervene in the evolution of any sociotechnical system where individuals can amplify content while referencing its originator.

Key words: social networks, triad transitivity, tertius iungens, social media

Significance Statement

When content goes viral on social media, or is amplified by a prominent account, the original poster of the content tends to gain followers. We introduce a process, attention brokerage, by which prominent individuals create new ties in social networks by amplifying others. When an attention broker amplifies another individual, people who follow the attention broker then follow the amplified individual. This provides insight into the processes at work that guide the emergence of network structure from within.

Introduction

Local causal processes influence the global structure of social networks; this can be seen in Schelling's work on segregation in 1960s U.S. neighborhoods [64] or Russo's work on radicalization on present-day Reddit [63]. In this work, we focus on an endogenous process by which open triads become transitive in a social network. Transitive triads tend to appear more often in networks than random chance and lower-level network characteristics (e.g. degree distribution, dyad census) would suggest [20]. Theories of social balance and tie strength can help explain the tendency social networks have towards forming transitive triads over so-called "forbidden" triads (see [21, 22, 40, 16]), but they do not provide a direct causal mechanism by which an open, "forbidden" triad becomes a transitive one. Discovering such a mechanism requires attention to, and measurement of, the small yet concrete changes that culminate in the emergent structures of attention we see on social media. Here, we outline such a mechanism that theoretically applies to any social network where users can amplify content with attribution to the content's author, such as citation graphs or corporate reporting hierarchies, and empirically demonstrate this mechanism on Twitter/X.

On social media specifically, influential users of various kinds guide and even gatekeep public discourse [39, 82, 74, 79, 55]. Online interactions and spreading processes on multiple social networking sites are driven in large part by intermediaries, and popular users have a significant effect on what those processes look like [2, 27, 28]. We examine the consequences a spreading process involving influential users can have for emergent properties of network structure. Specifically, we study two influential Twitter/X users who each regularly spread content related to a particular political agenda. We observe a pattern where followers of the influential users, whom we call *attention brokers,* follow accounts whose content the attention broker amplifies at a higher rate than the content's organic spread would otherwise indicate. Put simply, we empirically demonstrate how local, causal processes change the global structure of a social network, with attention brokers playing a central role.

We make two contributions – one theoretical and one empirical. Our first contribution is a theory of *attention brokerage.* Attention brokerage is the mechanism by which amplification of content by a specific high-degree node leads to the formation of transitive triads and a corresponding reallocation of attention within the network. Each of these transitive triads includes the attention broker, one of their followers, and the amplified content's author. The *attention broker* is an extension of Obstfeld's *tertius iungens*, or "third who joins" [57]; the attention broker, then, is the *tertius amplificans,* or "third who amplifies."

Our second contribution is empirical evidence of attention brokerage on Twitter/X, a social media site with directed (one-way) following ties and ample amplification affordances. Using a novel data collection method, we are able to identify the time that one Twitter/X user followed another within arbitrarily accurate bounds. This improves on prior approaches that use the temporal order, but not specific time, in which following events took place. With time-bounded following events, we are able to infer how many new followers of account amplified by an attention broker followed that account as a direct result of being amplified, accounting for background spread of the amplified content by using non-followers of the attention broker as "untreated" units. In short, we can infer amplification's causal effect on rates of transitive triad formation.

Background

Attention on Social Media

Attention patterns on social media have significant consequences. Political journalists' attention to insular bubbles might lead them to miss important stories [73]; users in polarized media environments stop paying attention to their cross-partisan contacts [72]; and algorithmic factors influence who gets attention, potentially silencing marginalized voices [42]. Wolf et al. (2022) found that the scope of topics that successive cohorts of new Twitter/X users paid attention to decreased over time, reflecting the increasingly overwhelming volume of content available on the platform.

At the same time, the ability to collectively generate and spread overwhelming volumes of content for purposes of mass mobilization has proven its significance many times over [28, 29, 82, 7, 78]. If follower count signals access to attentional resources and, therefore, is a simple measure of social capital, then possessing the ability to redirect attention in a lasting fashion implies an ability to redirect the allocation of social capital [48, 59, 75, 1]. Here, we define this *attentional capital* as the latent ability of an actor to draw attention to their content; when an actor amplifies content, they change the distribution of attention around them, reallocating it to the individual being amplified.

Social capital on Twitter/X is most visibly enacted in amplification processes. When many people collectively pay attention to a piece of content, its reach can be attributed to some combination of broadcasting (spread from one entity to many) and person-to-person spread (what we traditionally think of when we think of "virality") [27]. Twitter/X allows users to have millions of followers; a highly followed account therefore leverages massive online social capital, in the form of latent attention, when they broadcast content to their followers. Influential users of various kinds guide and even gatekeep public discourse on social media, and their political impact is still being discovered [39, 82, 74, 79, 55]. Person-to-person spread, in the form of the retweet or quote tweet, repeatedly leverages attentional capital, albeit at a smaller scale. There has been a great deal of research into what makes a particular tweet go viral in terms of its content; the literature tends to focus mainly on the impact of particular platform affordances or the network structure of viral spread [56, 37, 38, 27]. Of particular relevance for the present work, Goel et al. find that content diffusion largely depends on the "largest broadcast" – the highest degree individual that shares that content [27].

The structure of one's network is one aspect of social capital; social capital can therefore be understood to include easy access to a node with high social capital, or location in a structurally optimal local neighborhood [48]. Social capital is conventionally considered to be imposed by the network on the individual in some way. Burt, Bourdieu, and Granovetter all see social capital as a consequence of one's position in the network in some way, meaning that social capital is imposed by the network on the individual [12, 32, 8]. Putnam differs somewhat in defining social capital as the network's collective capacity to realize communal goals [59]. The attention broker reshapes the collective attention and, perhaps, *intention*, of the users around them, so attention brokers fit better into a theoretical landscape where Putnam's definition holds.

The ability to shape social capital is highly consequential to the distribution of social capital in a network. One possible mechanism for enacting such power is through amplification, where a user with a large platform introduces their followers to a novel account, directing attention to it. Most platforms' affordances for "friending" or "following" make the relationships permanent by default; Liang et al. found that only 2.89% of ties initially present in a sample of Twitter/X users disappeared over a four-month period [47]. Previous work has found that the retweet-follow sequence – amplification shaping social capital – has significant impact on Twitter/X [3]. Similarly, influencers' endorsement of products increases those products' credibility amongst their followers [14]. The fact that such processes are mediated by social media in some way is crucial; thanks to platform affordances, it requires very little effort to durably redirect one's attention, and personalized curation algorithms aggressively surface content generated by, or amplified by, accounts who have engaged the user in the past. A user who is capable of such rewiring does not necessarily need to have a large following in terms of absolute numbers. If they are able to engage their audience enough to alter their following patterns, then they possess sufficient capacity to change the distribution of attentional resources in their immediate surroundings.

Local Network Phenomena with Global Implications

Since attention is a directed phenomenon, the present work uses the *transitive triad* as its unit of analysis. A transitive triad involves *directed* relationships between three nodes, A, B, and C, such that (at minimum) $A \rightarrow B$, $B \rightarrow C$, and $A \rightarrow C$. Such triples of individuals are the smallest group size in which alliances can form, brokers can operate, and mediation can occur [68, 19]. For this reason, several theories relevant to the present work – most prominently Granovetter's strength of weak ties and Burt's structural holes [19, 10, 32] – use the triad as their basic unit of analysis. Directed social networks tend towards transitivity, and "open" or "forbidden" (i.e. intransitive) triads are unstable over time [40, 16, 20].

However, the exact mechanism (or set of mechanisms) by which triads become transitive over time is unclear. Individuals in the middle of triads can work to close them (or inhibit closure), but existing work is not able to empirically determine and validate the causal chain of events leading to closure (or prohibiting closure) [76, 36, 81, 9, 21]. Similarly, theories of brokerage, a process where a third party helps two other parties coordinate their actions, tend to emphasize that the broker spans a structural hole in the network and profits from their unique position [58, 68, 12]. They do not, however, examine the causal mechanisms leading to the broker profiting from their position in the network.

Small, locally controlled processes can change the global structure of a social network. Schelling provides a famous example in which individual agents are assigned to desire a certain fraction of their immediate neighborhood to be of their same type [64]. This local, endogenous decision-making process produces striking global segregation patterns even without exogenous intervention. More recent work on emergent network structure governed by local processes has studied phenomena like community self-organization and cultural tolerance using adapted Schelling models [5, 31]. Empirical work using social media data has found that economists have improved job market prospects when their job market papers are amplified by prominent scholars [60], and that users who receive replies from members of fringe subreddits are more likely to subsequently post on those subreddits [63]. In each of these cases, individual instances of amplification and attention were key in changing individual trajectories and altering the collective future of the network itself. In this work, therefore, we develop a causal theory of third-party amplification leading to triad transitivity and provide empirical evidence for the process' effect on link formation.

The **Tertius Iungens**

The *tertius iungens*, or "third who joins," is a strategic orientation towards creating ties between previously disconnected alters or helping connected alters coordinate with each other [57]. Obstfeld, who developed the idea of the *tertius iungens* in 2005, found that a *tertius iungens* orientation was correlated with innovation in the context of the firm. In the context of the firm, *tertius iungens* actors alter the shape of collaboration networks; on social media, an actor with a *tertius iungens* orientation employs affordances unique to social media to reshape the flow of attention, just as was described in the previous section. More recent work has found that the *tertius iungens* orientation positively impacts knowledge sharing behavior on enterprise social media and moderates the relationship between social network properties and innovation in biomedical research [45, 49].

Obstfeld argues that *tertius iungens* behavior is a form of brokerage; while other scholars have defined brokerage in ways that preclude or deemphasize such behavior, both Burt and Obstfeld use a more inclusive definition of brokerage [50, 30, 11, 57]. Exclusive definitions of brokerage tend to emphasize a different orientation: the *tertius gaudens* [68]. The *tertius gaudens* is the "third who enjoys;" this is the prototypical broker who spans a structural hole and increases their own social capital within the network by doing so [12, 68]. Such actors also increase the collective efficiency and knowledge of their entire network [32]. However, the *tertius gaudens* orientation is not appropriate for all networks; in some settings, such as public relations, a "power-hungry" *tertius gaudens* orientation may be suboptimal compared to collaboratively oriented *tertius iungens* behavior [44]. In other settings, however, maintaining a *tertius gaudens* orientation may not even be feasible.

Most of the work on brokerage and the *tertius iungens* orientation examines social networks in the context of firms and organizations; therefore, some of the concepts from the literature do not neatly translate to behavior on social media *per se*. For example, it would be difficult to be a *tertius gaudens* on Twitter/X due to the simple fact that establishing and maintaining following ties does not require deep investment on the part of either the follower and the followee. Unless a user has locked their Twitter/X profile, that profile is as accessible to a close friend of the user as it is to any other Twitter/X user. Controlling the flow of information between two users such that one can benefit from doing so is therefore infeasible for any aspiring Twitter/X *tertius gaudens*.

Obstfeld's *tertius iungens* necessarily operates at a small scale; coordinating or durably joining two individuals in an organizational context takes time and effort. If this joining occurred on a large scale, two ingredients would have to be present: large-scale amplification and low-effort, lasting tie formation. Twitter/X has made amplifying other users' content with attribution (i.e. retweets and quote tweets) simple and intuitive; most social media platforms offer similar or analogous affordances. Generally speaking, then, users are easily able to point their followers to accounts or content that they found interesting or informative and provide attribution to its creator. Additionally, as we have seen earlier in this section, most following ties tend to persist in the long term; they do not require sustained cultivation as would be required in the context of a company. If large-scale amplification can lead to following ties forming, something akin to the *tertius iungens* orientation is well positioned to thrive on social media platforms with such affordances.

Attention Brokers

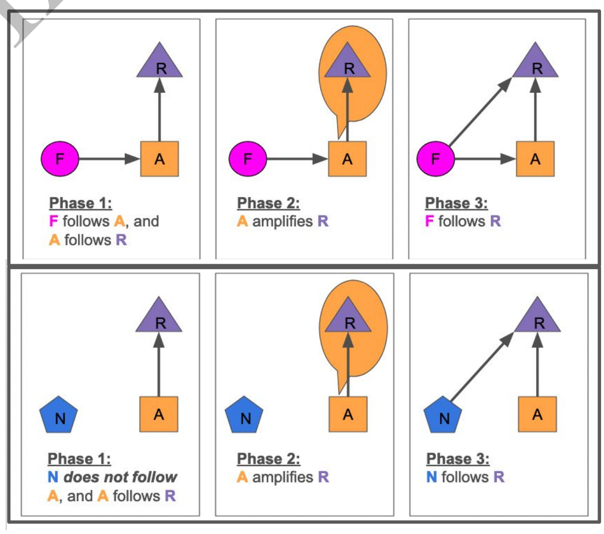

Fig. 1. Schematics of the attention brokerage process (top) and the "background virality" process (bottom).

We therefore focus on an extension of *tertius iungens* behavior wherein a well-established social media user (with a large number of attentive followers) facilitates ties between their followers and accounts they amplify. We introduce the attention broker, or *tertius amplificans*, the "third who amplifies," as an adaptation of the *tertius iungens* to a setting where amplification can occur at enormous scale. Since the *tertius amplificans*/attention broker is a niche that only exists in an environment with functionality for many-to-many amplification, it is a novel extension of the *tertius iungens*. It is well-established in Twitter/X users' understanding of the platform that retweets by more famous individuals lead to more followers, but the causal mechanism has not been empirically validated to date.

Specifically, attention broker behavior occurs when an individual A, the attention broker, **amplifies** content created by another individual R in a way that **attributes credit** to R. Individuals whose attention is directed to A, perhaps through following ties on social media (as in the present work), see that A has amplified R's content. A's followers, if interested in R, **then direct their attention in an enduring way towards** R, likely also in the form of new following ties. These new followers probably have latent tendencies that make them more likely to follow R; we are careful to note that attention brokerage is a process that relies on **exposure** rather than **persuasion**. The top line of Figure 1 provides a stylized example of how this might play out on social media. In the context of Twitter/X, and for the purposes of our empirical analysis, "amplification" refers to a simple retweet. At the start of the process, A's attention is directed at R, and A's followers' attention is directed at A, forming many "forbidden" triadic configurations [68]. The attention broker A resolves some number of these instabilities when they amplify R; an attentional triad consisting of A, R, and a follower becomes transitive when the follower follows R.

The ability to broadcast information to an audience at scale is central to the functionality of the attention broker. Obstfeld's empirical study linking *tertius iungens* behavior with innovation took place in the context of a 1000-person company with 440 professional-level employees eligible for his survey [57]. A prototypical *tertius iungens* interacts with individuals and forms ties on the individual level. The theory of the *tertius iungens* presupposes the ability to introduce two individuals and spend time cultivating the burgeoning connection. The attention broker, however, amplifies another user's content to their audience of thousands or millions. Their followers vet new ties and create following links on their own. Attention brokers do not actively persuade their followers to build ties; instead, they expose their followers to accounts that they may follow if the account matches a given follower's latent preferences. Most followers will not follow any one retweeted account, but the frequency and scale of amplification that broadcasting media affords for means that over time, an effective attention broker will have a sizeable effect on the shape of the corner of the social network they occupy. An impactful attention broker builds social capital in the form of audience trust and, in turn, reshapes the social/attentional capital around them. As they change the shape of attention in their vicinity, they alter the future of the network in small but measurable ways.

Data & Measurement

Case Selection

In work where the number of cases studied is, by necessity, small, the cases chosen can significantly impact the results of the research [26]. Here, we choose two accounts that differ from each other in several important ways despite both being attention brokers; namely, they have differing follower counts, espouse contrasting political views, are famous for different reasons, and one posts under their real name while the other plays a character. Our aim in choosing two very different cases is to illustrate how fundamental attention brokerage is to modern-day social networks equipped with broadcasting and amplification affordances. While the two accounts differ greatly, both have provably and persistently used their explicitly non-political social capital to permanently redirect their followers' attention to a political cause. By using two distinct cases, we ensure the results are not idiosyncratic to one account or the other.

Jorts the Cat's fame is specific to social media; users first became aware of him through a viral Reddit post [70] and Jorts' Twitter/X account quickly became a famous Twitter/X character. Jorts' pro-union politics were not clear at the outset, but he quickly became known as a supporter and amplifier of labor activism [80]. The account currently has around 200 thousand followers. We analyzed Jorts' attention broker behavior for the period of time between December 14, 2021, which is when the account was created, and March 13, 2022 (the day that the Jorts dataset was originally collected). J.K. Rowling, in contrast, is famous for reasons independent from social media; she is the billionaire author of the *Harry Potter* book series. Since 2019, Rowling has become an increasingly vitriolic TERF (trans-exclusionary radical feminist), and regularly amplifies transphobic speech to her 14 million followers [35, 62]. The last *Harry Potter* book was published in 2007, and Rowling's TERF politics became impossible to ignore in the late 2010s [24]. Rowling currently has over 14 million followers, an order of magnitude more than Jorts has. We analyzed her attention broker behavior between June 15, 2018 and March 1, 2023. Setting the start date in 2018 allows us to examine Rowling's behavior for a substantial stretch of time before she began espousing TERF ideology. The end date of March 1, 2023 reflects when the data collection occurred and encompasses a period of increasing TERF activism on Rowling's part.

Jorts' social media presence advocates for progressive and anti-hegemonic labor activism, while Rowling's TERFism is fundamentally regressive and perpetuates hegemonic, cissexist norms. Moreover, Rowling has an order of magnitude more followers than Jorts does, and her role in public discourse differs greatly from Jorts'. Both of these accounts do, however, exhibit impactful attention broker behavior. By using two very different cases, we hope to make it very clear that attention brokerage is something that transcends political agendas or specific platforms. Instead, it is a fundamental process in social networks with large-scale amplification and broadcasting affordances. This kind of large-scale network wiring caused by a single individual is not accounted for in mechanistic social network evolutionary models, which is consistent with the logic of these models [18, 17, 69, 34, 71, 6, 41, 51]. Exceptional individuals with influence beyond their immediate alters are not considered in typical modeling scenarios, yet attention broker behavior has real, measurable effects on the Twitter/X information ecosystem.

Measuring Time-Bounded Following Events

We use the Twitter/X API both to gather data about following behavior and to pinpoint amplification events by attention broker accounts. The Twitter/X V1 API supports multi-page queries through the use of a *cursor*, a long string of numbers that is returned along with a list result objects. The intended use case for the cursor is for users to send the most recently received cursor in their next query; the API then can pick up where it left off and return the results that directly follow the "page" of results that the user had most recently received. The cursor is also a modified Unix nanosecond timestamp (see [54] for more discussion of Twitter/X's use of Unix timestamps), and following events are returned by the V1 API in reverse chronological order.

Traditionally, scholars who are interested in the timing of following events are limited to triangulating timings based on some combination of the order of following events, account creation times, and repeated queries (see [3, 53] for examples). However, it was possible for a user of the V1 API to specify their own cursor, converted from a Unix timestamp, and only receive following events that occurred before that point in time [46]. We use this novel technique to only collect following events that occurred directly before and directly after accounts were amplified by the attention broker users. We note that this measurement of the precise moment of tie creation opens interesting possibilities in studying network evolution more generally when studying Twitter/X follower networks[1]. Any researchers who have collected and stored raw follower data have also incidentally collected information on when that edge was created. However, this method does **not** allow us to measure unfollowing events or follows by subsequently deleted accounts; only following ties that exist at the time of data collection can be recorded with this method. The scale of the data collection for this project precluded us collecting perfectly accurate time bounds, but the level of granularity we selected allows us to enumerate which follows occurred within the two-week periods we delineate directly before and after an attention broker's retweet.

Account Types

Both Jorts the Cat and J.K. Rowling are well known for supporting specific causes – union activism and labor rights for Jorts, and TERFism (trans-exclusionary radical feminism) for Rowling. We therefore compute effect size separately for accounts related and unrelated to each attention broker's cause. In Rowling's case, we categorize accounts by "interest actor" status in addition to labeling them as TERF or non-TERF. Interest actors are influential users who have a sizeable online following due to their "expertise, authority, or professional position," as a rough analogue to the calls to action studied by Goel et al. [55, 27]. We break out Jorts' brokerage effects into union-related accounts and non-union-related accounts. Labeling guidelines for both Rowling and Jorts' retweeted accounts can be found in the SI, in Tables S2 and S7. During the process of labeling accounts retweeted by Jorts as union-related or non-union-related, participation by a third researcher to resolve disagreements between the original two labelers was required 7.59% of the time. Similarly, Rowling's attention brokerage effects are broken out by the two binary labels of "Interest Actor" and "TERF". During labeling, participation by a third researcher to resolve disagreements between the original two labelers was required 5.1% of the time for the "TERF" label and 13.1% of the time for the "Interest Actor" label.

Data Collection & Processing

1. *Timeline collection of attention broker accounts*: Here, we use the focalevents Github package [23], which is a tool for organizing data collected with the Twitter/X V2 API, to collect Jorts' and J.K. Rowling's entire timelines. We then filter the timelines for all simple retweets (i.e. not quote tweets) by Jorts and J.K. Rowling. We avoid quote tweets because of a common practice, colloquially known as "dunking," where a user quote tweets a tweet they disagree with, adding commentary or insults as they do so. This gives us all of the instances when the attention broker accounts of interest amplified another account without "dunking" behavior.
2. *Time-bounded following event collection*: Using a custom Python script and the Twitter/X V1 API, we use the cursor parameter to collect following events between specified time bounds. Specifically, for each account that was retweeted by Jorts or J.K. Rowling during the period of interest, we collect all followers accumulated by an account in the two weeks directly before and directly after the first time that Jorts or J.K. Rowling retweeted the account. For Jorts, we look at attention broker activity between December 14, 2021, which is when the account was created, and March 13, 2022. We measure J.K. Rowling's attention brokerage between June 15, 2018 and March 1, 2023.
3. *Time-bounded attention broker follower collection*: Again, using a custom Python script and the Twitter/X V1 API, we collect pages of followers for both Jorts and J.K. Rowling. For each page of followers, we retain the cursor timestamps bounding the following events and map each follower to a period of time in which they must have followed Jorts or J.K. Rowling. We have time bounds for all Jorts' accumulated followers before mid-2022 and for Rowling's accumulated followers between June 15, 2018 and March 2023. Due to data limitations, we were not able to collect time bounds on all of Rowling's 300 million followers.
4. *Labeling Retweeted Accounts*: Four researchers labeled the accounts that were retweeted by either Jorts or J.K. Rowling in the periods of interest. The labeling guidelines, which can be found in the Appendix, are designed to flag ideological accounts (either TERFs for J.K. Rowling or labor activists for Jorts), as well as interest actors in Rowling's case [55]. There are at minimum two labelers assigned to each account; in cases of disagreement, a third labeler steps in to resolve the discrepancy. For Jorts The Cat, coders disagreed on 49 of the 646 accounts labeled; 112 of the accounts were not labeled because they were deleted, suspended, inactive, or protected. Coders disagreed on 27 of the 534 accounts retweeted by J.K. Rowling, and 77 of the accounts were not labeled for the same reasons enumerated previously. Tables with coding guidelines can be found in the Appendix.

Analysis

Background Virality & Non-Follower Controls

Attention brokers do not operate in a vacuum. A tweet retweeted by an attention broker may be in the throes of virality before the attention broker intervenes in the system, for example. This would cause an increase in follower accumulation around the time of the retweet, and we must therefore separate the effects of this viral spread from those of attention brokerage. We call the viral spread a tweet may or may not be undergoing around the time it is retweeted by an attention broker "background virality." Figure 1 contrasts attention brokerage (top) with background virality (bottom). In order to determine how much of the increase in following rate experienced by a

[1] For more information on the cursor trick, see `https://github.com/asmithh/jkr-follower-accumulation`

user is related to attention brokerage *only*, we measure how the rate of following by **followers** of the attention broker *and* that of **non-followers** changes after an attention broker's retweet. On average, the content seen by followers and non-followers of an attention broker who go on to follow the same account around the same time is not likely to differ substantially *except* in that followers are much more likely to be exposed to the attention broker's retweet. If exposure to the attention broker's retweet substantially changes the rate at which followers follow the retweeted account, then we can conclude that attention brokerage has an impact distinguishable from virality. Our inferential strategy closely follows that of McCabe et al. [52] in that we look at the differences in behavior before and after an inciting event between followers and non-followers of relevant accounts.

We therefore use the set of time-bounded following events and the time-bounded follower list for each attention broker to detect the formation of two kinds of motifs; the first is a transitive triad, and the second is an open triad. First, we flag all transitive triads consisting of a follower F, the attention broker A, and an account R retweeted by the attention broker that match the following criteria:

1. F follows A and F follows R.
2. The edge connecting F to A existed before A retweeted R.
3. The edge connecting F to R formed either in the two weeks before A retweeted R or in the two weeks directly after A retweeted R.

This comprises our "treatment" group; it is the set of following events that resulted from the combination of background virality and attention brokerage.

Second, we flag all open triads consisting of a non-follower N, the attention broker A, and an account R retweeted by the attention broker matching the following criteria:

1. N does *not* follow A prior to N following R and N follows R.
2. The edge connecting N to R formed either in the two weeks before A retweeted R or in the two weeks directly after A retweeted R.

Counting these open triads lets us determine the extent to which background virality *alone* increased the rate at which users followed R. In other words, the non-followers and their open triads comprise our "control" group. For each of the retweeted accounts R, we measure the number of times open triads and transitive triads that met the above conditions formed during the two-week periods before and after the attention broker A retweeted R. This forms the basis of our analysis.

Quantifying Attention Broker Behavior

We use difference-in-differences to infer attention brokers' direct effects on following ties by estimating the extent to which following an attention broker increases a user's tendency to follow an account after it is retweeted by that attention broker. For each account R retweeted by an attention broker, we examine the behavior of two sets of users, as detailed in the previous section: users that followed the attention broker at the time of the retweet ("followers"), and users that were not following the attention broker at that time ("non-followers"). We measure the daily estimated rates of accumulation of following ties to the retweeted account R by followers *and*

Table 1. Estimated Attentive Populations for Attention Brokers' Followers and Non-Followers

Population	**Estimated Size**	**Stderr**
Jorts, Followers	163987.4	223.4
Jorts, Non-Followers	17890822	10145.7
JKR, Followers	841164.4	1579.5
JKR, Non-Followers	2675853	100496.9

non-followers in the two weeks directly before and after the attention broker retweets R. For each chunk of time-bounded followers accumulated by an account and collected using our data collection method, we distribute the followers linearly over the indicated time bounds. This lets us estimate how many followers each retweeted account R accumulates each day. Some accounts that were followed only have three chunks of followers returned by the API (one measured on the day of the retweet, one two weeks after, and one two weeks before), while others have 50 or more in total. For Jorts the Cat, the average number of chunks obtained per account is 32.7; for J.K. Rowling, it is 4.4. The precision of the time bounds we have varies accordingly and likely contributes to uncertainty in the final analysis. We do not believe that the existence of small chunk sizes materially affect our inferences because these more coarse-grained measurements would bias our analysis towards null results. Histograms of the distribution of chunk counts can be found in the Supplementary Information in Figures S1-S2.

For both followers and non-followers, we divide the approximate daily follower accumulation rate by the estimated attentive population size. We estimate the number of accounts, either following the attention broker or not, who are attentive on the site and therefore are capable of noticing viral or retweeted content and following its author, using a mark-recapture population estimation algorithm from the free software package Project MARK [65, 66]. This algorithm assumes wildlife from a particular population are captured, uniquely marked, and then released (and potentially recaptured). It is possible to estimate the population size by aggregating individuals' capture histories over several marking periods. In this setting, a "capture" event for an individual account occurs when it follows an account that was retweeted by the attention broker on a specific day. We then obtain "capture histories" for all known followers and non-followers of the attention broker and use the POPAN implementation of the Jolly-Seber model to compute approximate population size [4, 43, 67]. The estimated population sizes, along with their standard errors, can be found in Table 1.

We now have a time series of estimated follower accumulation rates by followers and non-followers for each retweeted account R.

We specify the following model, using a two-way fixed effects formulation:

$$Y_{it} = \mu_i + \mu_t + \sum_{k=-14}^{-1} \tau^k D_{it}^k + \sum_{k=0}^{14} \tau^k D_{it}^k + \epsilon_{it} \quad (1)$$

Here, the variable μ_i refers to individual account-level fixed effects, and μ_t to temporal fixed effects. t refers to the time step, where 0 is the day on which the attention broker's retweet occurred. D_{it}^k is a treatment lead/lag variable, and τ^k is the treatment effect on day k. See Figures 2 and 3 for a visualization of the average pre-trends and post-trends for followers and non-followers (with standard deviations shaded in a lighter color).

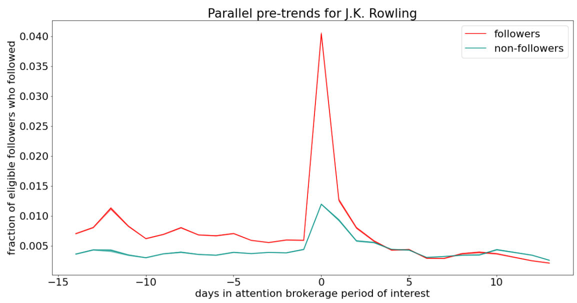

Fig. 2. Pre-trends plot for J.K. Rowling

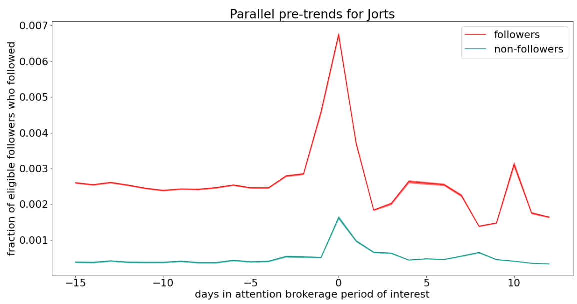

Fig. 3. Pre-trends plot for Jorts The Cat

While this shows some evidence of pre-trends consistent with followers being more likely than non-followers to follow the amplified account prior to amplification, we use the sensitivity analysis described in Rambchan and Roth [61] to characterize the extent to which the effects we identify are robust to violations of the parallel trends assumption. The results of these analyses can be found in Tables S5 & S6 (for Jorts) and S12-S15 (for Rowling). The parallel trends assumption is necessary for differences-in-differences analyses; it assumes that if the treatment group (i.e. an attention broker's followers) were not exposed to "treatment," the average change they experienced would be the same as that of the control group after the time at which intervention would have occurred. The implications for the robustness of our results, given these analyses, will be discussed in the Results section. We run two-stage differences-in-differences event study analyses for each attention broker and each account type, along with the aforementioned sensitivity analysis [13, 25, 61]. We assess the coefficients for treatment effect size on each account type, as well as their robustness under different magnitudes of violation of the parallel trends assumption, in the **Results** section. Additionally, we note that the coefficients obtained here are **conditional average treatment effects**; we make no claims that any one individual-level outcome is the result of attention brokerage. Instead, we obtain the increase in following rates **in expectation** due to attention brokerage. A specific individual following event could be the result of another attention broker's intervention, an algorithmic intervention, or a recommendation from another social networking site, for example. Our analysis indicates the extent to which Jorts and J.K. Rowling's retweets are expected to increase the rate at which retweeted accounts accumulate followers.

Results

Jorts The Cat

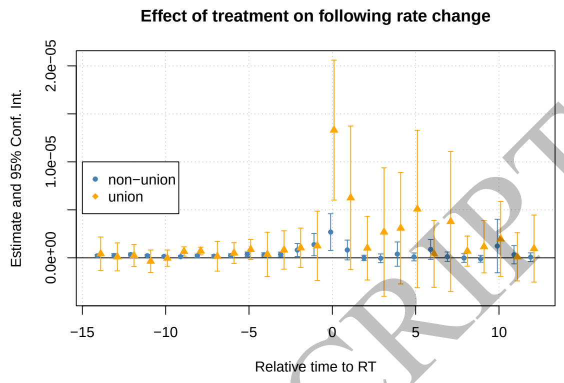

Fig. 4. Jorts The Cat: Two-Stage Differences-in-Differences Effect Sizes Over Time with Standard Errors by Pro-Union Status

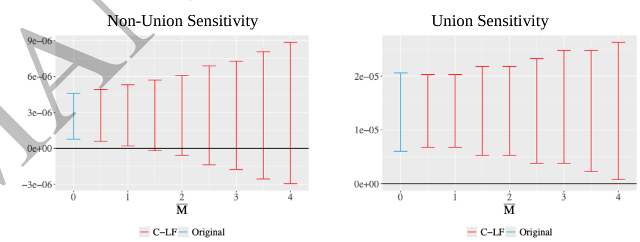

Fig. 5. Jorts The Cat: Sensitivity Analysis by Account Type

Jorts The Cat exhibits impactful attention broker behavior on Twitter/X. Jorts began his life on Twitter/X as a humorous cat account and quickly began incorporating pro-union content into his posts. For both union-related accounts and non-union-related accounts, the effect size of "treatment" (i.e. Jorts' retweet) on the day of Jorts' retweet was significant. The fact that effect sizes are also positive and significant prior to Jorts' retweet, but are also much smaller than on the day of Jorts' retweet, suggests that incidental prior exposure also plays a role in follower accumulation. Followers of Jorts who have a latent predisposition to follow accounts that Jorts might retweet could be exposed to such an account prior to Jorts' retweet and choose to follow it at that (earlier) point in time. The exact results of the analysis can be found in Tables S3 and S4.

However, a sensitivity analysis using methods from Rambachan and Roth (see Figure 5) indicates that our results are likely robust to violations of the parallel trends assumption [61]. Rambachan and Roth's method allows us to quantify how many times larger the post-intervention violation of the parallel trends assumption must be, compared to the magnitude of any pre-intervention violations, for our results to become invalid. Our sensitivity analysis assesses different values of $\bar{M}$, where the post-retweet violation of parallel trends is $\bar{M}$ times larger in magnitude than the maximal value pre-retweet. While the threshold at which significance becomes questionable for non-union accounts is lower than for union accounts, the violation of the parallel trends assumption post-retweet would have to be greater than 4 times larger than the maximum pre-retweet parallel trends violation in order for Jorts' effect on union accounts to not hold.

J.K. Rowling

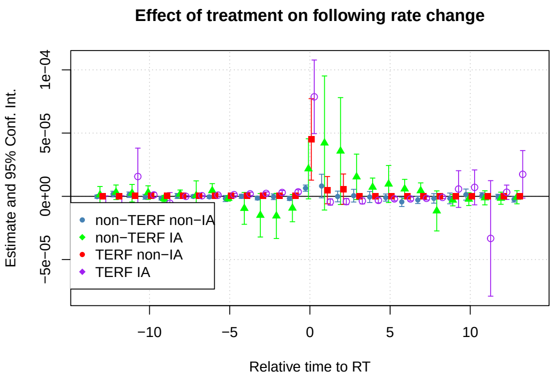

Fig. 6. J.K. Rowling: Two-Stage Differences-in-Differences Effect Sizes Over Time with Standard Errors by TERF and Interest Actor Status

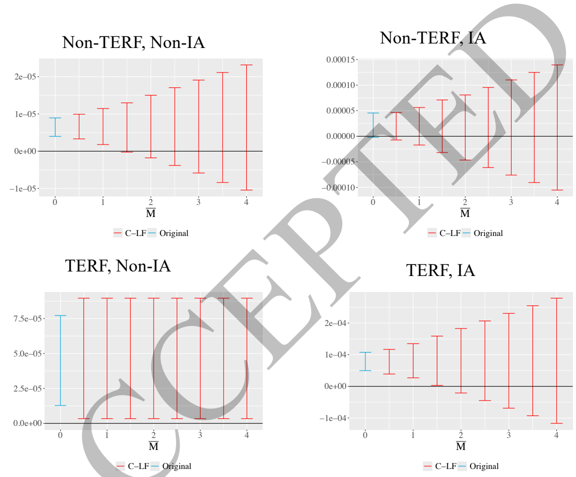

Fig. 7. J.K. Rowling: Sensitivity Analysis by Account Type

J.K. Rowling's attention broker behavior was significant for all account types on day 0, the day of the amplification event. Like Jorts The Cat, Rowling's impact as an attention broker is heterogeneous across account types. The "treatment effect" of Rowling's retweet on day 0 is significant for all account types, but it is smallest for non-TERF non-interest actor accounts. Her attention brokerage on day 0 for TERF accounts is significant regardless of prominence (i.e. interest actor status), although the effect size on day 0 appears to be largest for TERF interest actor accounts. The exact results of the analysis can be found in Tables S8-S11.

Additionally, the parallel pre-trends plot (Figure 2 shows that the following rate of Rowling's followers to the retweeted account R decreases after an initial spike post-retweet. This is corroborated by the event study plot (Figure 6, where coefficients for TERF interest actor accounts are slightly negative after the initial large spike. We believe that, similar to Jorts' retweets, Rowling's retweets **accelerate** the rate at which retweeted accounts follow accounts they were likely to follow anyway. This acceleration diminishes the pool of potential followers and leads to a lower subsequent rate in following several days after Rowling's retweet. According to our sensitivity analysis (see Figure 7), which is described in the previous subsection, Rowling's results hold for at least one account type even if the violation of the parallel trends assumption post-retweet is greater than 4 times than the maximum pre-retweet parallel trends violation.

Discussion & Conclusion

Implications

Due to the fundamental problem of causal inference, we cannot attribute any specific, individual following event to the actions of either attention broker. However, our analyses demonstrate that both Jorts and J.K. Rowling affect the shape and structure of their social networks through their amplification of like-minded accounts. For at least one account type, each attention broker was **impactful**, in that they accelerated the rate at which their followers followed accounts they retweeted **after the retweeting event** more than for non-followers. In other words, the increase in the rate at which "forbidden" triads transformed to stable, transitive ones due to attention brokerage and background virality was larger than the increase in the rate at which new forbidden triads formed due to background virality: attention brokers' intervention pushes the network to evolve. Future work, perhaps qualitative in nature, could seek to understand *why* attention brokers choose to amplify the accounts they do and, in line with DeVito et al. [15], elicit folk theorizations of attention brokerage among attention brokers and their followers. Understanding how these influential users and their followers make sense of the ways in which they are driving the evolution of their networked neighborhoods would be a valuable contribution to the CSCW (computer-supported cooperative work) literature.

We can conclude that transitivity is an emergent consequence of attributed resharing processes in social networks with directed links, by way of our proposed mechanism. Additionally, it is important to note with high likelihood that attention brokerage is occurring outside the two case studies presented in this article. While our analyses focus on two individual attention brokers, we emphasize here that the phenomenon is not limited to J.K. Rowling and Jorts the Cat. At scale, it is an important mechanism by which networks of attention evolve. One agent briefly shines a light on another, resulting in an incremental but long-standing increase in attention for the latter When an attention broker shares content by someone they are linked to, they indicate to their audience that that individual creates content that is worth paying attention to. Therefore, members of the attention broker's audience are more likely to form a link to an individual at the time when the attention broker amplifies that individual's content. Resharing processes like retweeting on Twitter/X, citing academic articles in one's own work, creating duets on TikTok, or circulating memos in corporate settings will naturally create transitive triads *and* stabilize the network. Indeed, the work of the attention broker offer a causal explanation for the prevalence of transitivity in social networks.

In a setting where the number of entities one can pay attention to far outstrips the human capacity for attention (i.e. modern social media), individuals must find viable ways to allocate their attention. One such strategy relies on trusted entities to indicate which individuals are worth paying attention to. This played out in an experiment run by Qiu et al., in which amplification by a prominent economist significantly increased candidates' success on the academic job market [60]. When amplification involves little friction on the part of the amplifier *and* amplification occurs with attribution, as is the case not only on social media but also in academic citation networks or modern-day news media, it is likely that attention brokerage is shaping the global network structure. This is especially true in situations where attention is "sticky" (such as following affordances on social media), as "stickiness" may also make attention brokerage itself a self-reinforcing process. If following one account that posts about labor activism leads a user to follow others, those novel following ties to labor activist accounts represent new opportunities for attention brokerage to reallocate that user's attention. Future work could explore the effects of **follower** ideology, or other forms of heterogeneity in attention brokers' audiences, on the efficacy of attention brokerage and, perhaps, explore how follower ideology and information diets change over time due to attention brokerage. It is entirely plausible that amplification facilitates a discovery process of users that share content to the liking of a potential recipient. For example, let us say that A follows B because they find content that B shares. B retweets content from C that A likes, and thus A subsequently follows C. Such processes could, over time, create links that join communities around shared interests, gradually producing "curation bubbles" where users circulate and view content that aligns with their worldview, regardless of source [33].

As questions of the societal impacts of social media become more and more salient, understanding how users' attention is (re)directed toward novel content creators is key. Social capital, and its evolution alongside the rise of the Internet, is well-studied in the literature, but the ability for individuals to reshape social capital in their vicinity, inasmuch as audience size equates to social capital online, has been less thoroughly examined. Social media may be the venue in which it is easiest to empirically measure attention broker behavior, but attention brokerage is not limited to social media. Put simply, in social networks where attention is sticky, the universe of entities to pay attention to far exceeds human attentional capacity, and amplification involves attribution to an original source, attention brokerage is likely to be at work shaping global network structure. This transcends social media and has implications for our broader understanding of social networks. As discussed in the *Background* section, attention brokers make the most sense in a world where social/attentional capital is a collective project [48]. This differs from definitions of social capital that define it as a factor that gives individuals an edge in a competitive landscape [12, 32, 8]. The attention broker uses their individual social capital to promote solidarity and cohesion, changing the shape of their network neighborhood's collective intent towards the agendas they amplify.

Limitations

Limitations of this study tend to stem from Twitter/X API rate limits and outages. For example, it was infeasible to collect the entire following network pertaining to the attention broker accounts and their followers; such a dataset, coupled with all the retweets of all accounts followed by all the attention broker accounts' followers, would allow us to make more fine-grained causal claims about following events. Scale limitations, particularly those imposed by the Twitter/X API's rate limits, also mean that this study was restricted to two cases. Algorithmic curation of users' timelines and recommended follows may also have affected our results; the fact that we observe a significant effect for multiple cases, and over a long period of time, indicates that the algorithm's existence does not override our findings *and* does not affect the attention brokerage process enough to render it invisible. The scale of the data we analyzed means that it was impossible to check whether all accounts were not bots; however, we believe that bots would be no more likely to appear in the "treatment" groups (i.e. followers of an attention broker) than in the "control" groups (i.e. non-followers).

Data Availability Statement

We have deposited fully anonymized data at SOMAR (the Social Media Archive, hosted at the Inter-university Consortium for Political and Social Research at the University of Michigan); it is currently ready for publication and has a DOI index of `https://doi.org/10.3886/3swn-td91`. The dataset is restricted and users must agree that they will not attempt to reidentify accounts in the dataset. The full dataset consists of the time-bounded follower counts, and their associated timings, for each retweeted account and attention broker followers and non-followers. All retweeted accounts' usernames and followers' user IDs are hashed using a non-reversible hash function for privacy. Additionally, we have uploaded a complete codebase, with documentation, to GitHub; it is currently private but will be made public upon this article's publication. This includes code for data collection, parsing, and analysis.

Acknowledgments

Dr. Pranav Goel provided comments on multiple drafts of this article, for which we are very grateful. Hana Chalmers, Samantha Furey, Sasheen Joseph, and Alexandra Madaras, all Northeastern University undergraduate research assistants, provided invaluable help by hand-labeling over a thousand Twitter accounts. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1938052. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Finally, we would like to thank the two anonymous reviewers who took the time to thoroughly and constructively review this manuscript. Their comments added substantially to the depth and impact of this article.

Competing Interest Statment

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1938052. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. Paul S. Adler and Seok-Woo Kwon. Social Capital: Prospects for a New Concept. *Academy of Management Review*, 27(1):17–40, January 2002. Publisher: Academy of Management.
2. Huda Alhazmi, Swapna S. Gokhale, and Derek Doran. Understanding social effects in online networks. In *2015 International Conference on Computing, Networking and Communications (ICNC)*, pages 863–868, February 2015.
3. Demetris Antoniades and Constantine Dovrolis. Co-evolutionary dynamics in social networks: a case study of Twitter. *Computational Social Networks*, 2(1):14, July 2015.
4. A Neil Arnason and Carl J Schwarz. Using popan-5 to analyse banding data. *Bird study*, 46(sup1):S157–S168, 1999.
5. V. Avetisov, A. Gorsky, S. Maslov, S. Nechaev, and O. Valba. Phase transitions in social networks inspired by the Schelling model. *Physical Review E*, 98(3):032308, September 2018. Publisher: American Physical Society.
6. Ginestra Bianconi, Richard K. Darst, Jacopo Iacovacci, and Santo Fortunato. Triadic closure as a basic generating mechanism of communities in complex networks. *Physical Review E*, 90(4):042806, October 2014. arXiv:1407.1664 [physics].
7. Leticia Bode. Facebooking It to the Polls: A Study in Online Social Networking and Political Behavior. *Journal of Information Technology & Politics*, 9(4):352–369, October 2012. Publisher: Routledge _eprint: https://doi.org/10.1080/19331681.2012.709045.
8. Pierre Bourdieu. THE FORMS OF CAPITAL. In *Handbook of Theory and Research for the Sociology of Education*, pages 241–258. Greenwood, Westport, C, 1986.
9. Michael Brzozowski and Daniel Romero. Who Should I Follow? Recommending People in Directed Social Networks. *Proceedings of the International AAAI Conference on Web and Social Media*, 5(1):458–461, 2011. Number: 1.
10. Ronald S. Burt. *Structural Holes: The Social Structure of Competition*. Harvard University Press, 1992.
11. Ronald S. Burt. The Contingent Value of Social Capital. *Administrative Science Quarterly*, 42(2):339–365, 1997. Publisher: [Sage Publications, Inc., Johnson Graduate School of Management, Cornell University].
12. Ronald S. Burt. Structural Holes versus Network Closure as Social Capital. In *Social Capital*. Routledge, 2001. Num Pages: 26.
13. Kyle Butts. did2s: Two-Stage Difference-in-Differences Following Gardner (2021), 2021.
14. Siyoung Chung and Hichang Cho. Fostering Parasocial Relationships with Celebrities on Social Media: Implications for Celebrity Endorsement. *Psychology & Marketing*, 34(4):481–495, 2017. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mar.21001.
15. Michael A DeVito, Darren Gergle, and Jeremy Birnholtz. " algorithms ruin everything" # riptwitter, folk theories, and resistance to algorithmic change in social media. In *Proceedings of the 2017 CHI conference on human factors in computing systems*, pages 3163–3174, 2017.
16. Seyedehmina Doroud. *Triads and micro-macro connections in online social networks - ProQuest*. PhD thesis, University of California Davis, 2014.
17. Thomas J. Fararo and John Skvoretz. E-State Structuralism: A Theoretical Method. *American Sociological Review*, 51(5):591–602, 1986. Publisher: [American Sociological Association, Sage Publications, Inc.].
18. Thomas J. Fararo, John Skvoretz, and Kenji Kosaka. Advances in E-state structuralism: further studies in dominance structure formation. *Social Networks*, 16(3):233–265, July 1994.
19. Katherine Faust. 7. Very Local Structure in Social Networks. *Sociological Methodology*, 37(1):209–256, August 2007. Publisher: SAGE Publications Inc.
20. Katherine Faust. A puzzle concerning triads in social networks: Graph constraints and the triad census. *Social Networks*, 32(3):221–233, July 2010.
21. Scott L. Feld. The Focused Organization of Social Ties. *American Journal of Sociology*, 86(5):1015–1035, 1981. Publisher: University of Chicago Press.
22. Scott L. Feld and Richard Elmore. Patterns of Sociometric Choices: Transitivity Reconsidered. *Social Psychology Quarterly*, 45(2):77, June 1982.
23. Ryan Gallagher. Social Media Focal Events, February 2023. original-date: 2021-08-17T19:51:50Z.
24. Abby Gardner. A Complete Breakdown of the J.K. Rowling Transgender-Comments Controversy | Glamour, April 2023.
25. John Gardner. Two-stage differences in differences. *Papers*, July 2022. Number: 2207.05943 Publisher: arXiv.org.
26. Barbara Geddes. How the Cases You Choose Affect the Answers You Get: Selection Bias in Comparative Politics. *Political Analysis*, 2:131–150, 1990. Publisher: Cambridge University Press.
27. Sharad Goel, Ashton Anderson, Jake Hofman, and Duncan J. Watts. The Structural Virality of Online Diffusion. *Management Science*, 62(1):180–196, January 2016.
28. Sandra González-Bailón, Javier Borge-Holthoefer, and Yamir Moreno. Broadcasters and Hidden Influentials in Online Protest Diffusion. *American Behavioral Scientist*, 57(7):943–965, July 2013. Publisher: SAGE Publications Inc.
29. Sandra González-Bailón and Ning Wang. Networked discontent: The anatomy of protest campaigns in social media. *Social Networks*, 44:95–104, January 2016.
30. Roger V. Gould and Roberto M. Fernandez. Structures of Mediation: A Formal Approach to Brokerage in Transaction Networks. *Sociological Methodology*, 19:89–126, 1989. Publisher: [American Sociological Association, Wiley, Sage Publications, Inc.].
31. C. Gracia-Lázaro, L. M. Floría, and Y. Moreno. Selective advantage of tolerant cultural traits in the Axelrod-Schelling model. *Physical Review E*, 83(5):056103, May 2011. Publisher: American Physical Society.
32. Mark S. Granovetter. The Strength of Weak Ties. *American Journal of Sociology*, 78(6):1360–1380, 1973. Publisher: University of Chicago Press.
33. Jonathan Green, Stefan D McCabe, Sarah Shugars, Hanyu Chwe, Luke Horgan, Shuyeng Cao, and David M J Lazer. Curation Bubbles. *American Political Science Review*, Accepted 2024, July 2024.
34. Yupeng Gu, Yizhou Sun, and Jianxi Gao. The Co-Evolution Model for Social Network Evolving and Opinion Migration.

In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, pages 175–184, New York, NY, USA, August 2017. Association for Computing Machinery.

35. Gina Gwenffrewi. J. K. Rowling and the Echo Chamber of Secrets. *TSQ: Transgender Studies Quarterly*, 9(3):507–516, August 2022.
36. Nir Halevy, Eliran Halali, and Julian J. Zlatev. Brokerage and Brokering: An Integrative Review and Organizing Framework for Third Party Influence. *Academy of Management Annals*, 13(1):215–239, January 2019.
37. Lars Kai Hansen, Adam Arvidsson, Finn Aarup Nielsen, Elanor Colleoni, and Michael Etter. Good Friends, Bad News - Affect and Virality in Twitter. In James J. Park, Laurence T. Yang, and Changhoon Lee, editors, *Future Information Technology*, volume 185, pages 34–43. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011. Series Title: Communications in Computer and Information Science.
38. Irina Heimbach, Benjamin Schiller, Thorsten Strufe, and Oliver Hinz. Content Virality on Online Social Networks: Empirical Evidence from Twitter, Facebook, and Google+ on German News Websites. In *Proceedings of the 26th ACM Conference on Hypertext & Social Media - HT '15*, pages 39–47, Guzelyurt, Northern Cyprus, 2015. ACM Press.
39. Jeff Hemsley. Followers Retweet! The Influence of Middle-Level Gatekeepers on the Spread of Political Information on Twitter. *Policy & Internet*, 11(3):280–304, 2019. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/poi3.202.
40. Paul W. Holland and Samuel Leinhardt. Transitivity in Structural Models of Small Groups. *Comparative Group Studies*, 2(2):107–124, May 1971.
41. Petter Holme and Beom Jun Kim. Growing scale-free networks with tunable clustering. *Physical Review E*, 65(2):026107, January 2002. Publisher: American Physical Society.
42. Benjamin N Jacobsen. Regimes of recognition on algorithmic media. *New Media & Society*, page 14614448211053555, October 2021. Publisher: SAGE Publications.
43. George M Jolly. Explicit estimates from capture-recapture data with both death and immigration-stochastic model. *Biometrika*, 52(1/2):225–247, 1965.
44. Michael L. Kent, Erich J. Sommerfeldt, and Adam J. Saffer. Social networks, power, and public relations: Tertius Iungens as a cocreational approach to studying relationship networks. *Public Relations Review*, 42(1):91–100, March 2016.
45. Kee-Young Kwahk and Do-Hyung Park. The effects of network sharing on knowledge-sharing activities and job performance in enterprise social media environments. *Computers in Human Behavior*, 55:826–839, February 2016.
46. Bob Leggitt. How To Find Out WHEN Someone Followed You on Twitter, November 2019.
47. Hai Liang and King-wa Fu. Information Overload, Similarity, and Redundancy: Unsubscribing Information Sources on Twitter. *Journal of Computer-Mediated Communication*, 22(1):1–17, January 2017.
48. Nan Lin. *Social Capital: A Theory of Social Structure and Action*. Cambridge University Press, May 2002.
49. Oscar Llopis, Pablo D'Este, and Adrián A. Díaz-Faes. Connecting others: Does a tertius iungens orientation shape the relationship between research networks and innovation? *Research Policy*, 50(4):104175, May 2021.
50. Peter Marsden. Brokerage behavior in restricted exchange networks. *Social structure and network analysis*, 7(4):341–410, 1982.
51. Matteo Marsili, Fernando Vega-Redondo, and František Slanina. The rise and fall of a networked society: A formal model. *Proceedings of the National Academy of Sciences*, 101(6):1439–1442, February 2004. Publisher: Proceedings of the National Academy of Sciences.
52. Stefan D. McCabe, Diogo Ferrari, Jon Green, David M. J. Lazer, and Kevin M. Esterling. Post-January 6th deplatforming reduced the reach of misinformation on Twitter. *Nature*, 630(8015):132–140, June 2024. Publisher: Nature Publishing Group.
53. Brendan Meeder, Brian Karrer, Amin Sayedi, R. Ravi, Christian Borgs, and Jennifer Chayes. We know who you followed last summer: inferring social link creation times in twitter. In *Proceedings of the 20th international conference on World wide web*, WWW '11, pages 517–526, New York, NY, USA, March 2011. Association for Computing Machinery.
54. Fred Morstatter, Harsh Dani, Justin Sampson, and Huan Liu. Can One Tamper with the Sample API?: Toward Neutralizing Bias from Spam and Bot Content. In *Proceedings of the 25th International Conference Companion on World Wide Web - WWW '16 Companion*, pages 81–82, Montreal, Quebec, Canada, 2016. ACM Press.
55. Laura Moses. Conceptualizing and Identifying "Interest Actors". March 2023. Publisher: Open Science Framework.
56. Krishnadas Nanath and Geethu Joy. Leveraging Twitter data to analyze the virality of Covid-19 tweets: a text mining approach. *Behaviour & Information Technology*, 42(2):196–214, January 2023.
57. David Obstfeld. Social Networks, the Tertius Iungens Orientation, and Involvement in Innovation. *Administrative Science Quarterly*, 50(1):100–130, 2005. Publisher: [Sage Publications, Inc., Johnson Graduate School of Management, Cornell University].
58. David Obstfeld, Stephen P. Borgatti, and Jason Davis. Brokerage as a Process: Decoupling Third Party Action from Social Network Structure. In Daniel J. Brass, Giuseppe (Joe) Labianca, Ajay Mehra, Daniel S. Halgin, and Stephen P. Borgatti, editors, *Research in the Sociology of Organizations*, volume 40, pages 135–159. Emerald Group Publishing Limited, July 2014.
59. Robert D. Putnam. *Bowling Alone: The Collapse and Revival of American Community*. Simon and Schuster, 2000. Google-Books-ID: rd2ibodep7UC.
60. Jingyi Qiu, Yan Chen, Alain Cohn, and Alvin E. Roth. Social Media and Job Market Success: A Field Experiment on Twitter. May 2024.
61. Ashesh Rambachan and Jonathan Roth. A More Credible Approach to Parallel Trends. *Review of Economic Studies*, 90(5):2555–2591, September 2023.
62. Hannah Ravell. #RIPJKRowling: A tale of a fandom, Twitter and a haunting author who refuses to die. *Public Relations Inquiry*, 12(3):239–270, September 2023. Publisher: SAGE Publications.
63. Giuseppe Russo, Manoel Horta Ribeiro, and Robert West. Stranger Danger! Cross-Community Interactions with Fringe Users Increase the Growth of Fringe Communities on Reddit, October 2023. arXiv:2310.12186 [cs].

64. Thomas C. Schelling. Models of Segregation. *The American Economic Review*, 59(2):488–493, 1969. Publisher: American Economic Association.
65. Carl James Schwarz and A Neil Arnason. A general methodology for the analysis of capture-recapture experiments in open populations. *Biometrics*, pages 860–873, 1996.
66. CJ Schwarz and AN Arnason. Chapter 12: Jolly-seber models in mark. *Program MARK: a gentle introduction*, pages 1–51, 2018.
67. GAF Seber. A note on the multiple-recapture census. biometrika 52: 249–259.. 1982. *Estimation of animal abundance and related parameters. Macmillan, New York, New York, USA*, 1965.
68. Georg Simmel. The Number of Members as Determining the Sociological form of the Group. II. *American Journal of Sociology*, 8(2):158–196, September 1902.
69. John Skvoretz, Katherine Faust, and Thomas J. Fararo. Social structure, networks, and E-state structuralism models. *The Journal of Mathematical Sociology*, 21(1-2):57–76, April 1996. Publisher: Routledge _eprint: https://doi.org/10.1080/0022250X.1996.9990174.
70. throwawayorangecat. AITA for "perpetuating ethnic stereotypes" about Jorts?, December 2021.
71. Riitta Toivonen, Lauri Kovanen, Mikko Kivelä, Jukka-Pekka Onnela, Jari Saramäki, and Kimmo Kaski. A comparative study of social network models: Network evolution models and nodal attribute models. *Social Networks*, 31(4):240–254, October 2009.
72. Christopher K. Tokita, Andrew M. Guess, and Corina E. Tarnita. Polarized information ecosystems can reorganize social networks via information cascades. *Proceedings of the National Academy of Sciences*, 118(50):e2102147118, December 2021. Publisher: Proceedings of the National Academy of Sciences.
73. Nikki Usher and Yee Man Margaret Ng. Sharing Knowledge and "Microbubbles": Epistemic Communities and Insularity in US Political Journalism. *Social Media + Society*, 6(2):2056305120926639, April 2020. Publisher: SAGE Publications Ltd.
74. Seth B. Warner. Analyzing Attention to Scandal on Twitter: Elites Sell What Supporters Buy. *Political Research Quarterly*, 76(2):841–850, June 2023.
75. Dmitri Williams. On and off the 'Net: Scales for Social Capital in an Online Era. *Journal of Computer-Mediated Communication*, 11(2):593–628, January 2006.
76. Michael Windzio. Social exchange and integration into visits-at-home networks: Effects of third-party intervention and residential segregation on boundary-crossing. *Rationality and Society*, 30(4):491–513, November 2018.
77. Frederik Wolf, Sune Lehmann, and Philipp Lorenz-Spreen. Successive Cohorts of Twitter Users Show Increasing Activity and Shrinking Content Horizons. *Journal of Quantitative Description: Digital Media*, 2, July 2022.
78. Masahiro Yamamoto, Matthew J. Kushin, and Francis Dalisay. Social media and mobiles as political mobilization forces for young adults: Examining the moderating role of online political expression in political participation. *New Media & Society*, 17:880–898, 2015. Place: US Publisher: Sage Publications.
79. Shuzhe Yang, Anabel Quan-Haase, and Kai Rannenberg. The changing public sphere on Twitter: Network structure, elites and topics of the #righttobeforgotten. *New Media & Society*, 19(12):1983–2002, December 2017.
80. Waiyee Yip. How a fluffy orange cat named Jorts stole the internet's heart and became the pro-labor icon 2022 didn't know it needed, 2022.
81. Pavel I. Zhelyazkov. Interactions and Interests: Collaboration Outcomes, Competitive Concerns, and the Limits to Triadic Closure. *Administrative Science Quarterly*, 63(1):210–247, March 2018.
82. Mathilda Åkerlund. The importance of influential users in (re)producing Swedish far-right discourse on Twitter. *European Journal of Communication*, 35(6):613–628, December 2020.